

Predicting Electrocardiogram and Arterial Blood Pressure Waveforms with Different Echo State Network Architectures

Allan Fong, MS^{1,3}, Ranjeev Mittu, MS², Raj Ratwani, PhD³, James Reggia, MD, PhD¹

¹University of Maryland, College Park, MD; ²Naval Research Laboratory, Washington DC; ³National Center for Human Factors in Healthcare, Washington DC

Abstract

Alarm fatigue caused by false alarms and alerts is an extremely important issue for the medical staff in Intensive Care Units. The ability to predict electrocardiogram and arterial blood pressure waveforms can potentially help the staff and hospital systems better classify a patient's waveforms and subsequent alarms. This paper explores the use of Echo State Networks, a specific type of neural network for mining, understanding, and predicting electrocardiogram and arterial blood pressure waveforms. Several network architectures are designed and evaluated. The results show the utility of these echo state networks, particularly ones with larger integrated reservoirs, for predicting electrocardiogram waveforms and the adaptability of such models across individuals. The work presented here offers a unique approach for understanding and predicting a patient's waveforms in order to potentially improve alarm generation. We conclude with a brief discussion of future extensions of this research.

1. Introduction

Intensive Care Units (ICUs) are designed to handle some of the most physiologically fragile patients in the hospital. As a result, ICUs utilize a wide spectrum of machines, technologies, and tests to help medical staff better understand and care for patients. However, the wide array of stand-alone machines often collect data and produce alarms and alerts independently, leaving the difficult integration tasks for the medical staff¹. Time sensitive decisions, including identifying non-critical alarms, are just some of the problems faced by ICU medical staffs. Studies have shown that staffs in ICUs face an extraordinary number of alarms each day, some as many as 1,000 alarms a day, many of which are non-actionable or not necessary for patient care^{2,3}. Excess amounts of non-critical alarms can lead to alarm fatigue which can adversely affect patient care^{4,6}. While there is ongoing research to effectively minimize false alarms, such as allowing nurses to adjust alarm thresholds, much work is still needed to improve classification techniques and systems in order to reduce false alarms^{5,7}. A remaining challenge is to develop algorithms robust enough to understand, integrate, and predict multiple physiological waveforms from patient data to better classify and interpret alarms.

The purpose of this paper is to explore the value of using of Echo State Networks (ESN), a particular type of recurrent neural network, to predict an individual's waveforms. Being able to forecast an individual's waveform could potentially offer much more information to the medical staff in addition to alarm classifications. ESNs were chosen because of their ability to accurately predict chaotic time series⁸⁻¹¹. In our research, these networks were trained to predict an individual's electrocardiogram (ECG) and arterial blood pressure (ABP) waveform data, which can potentially help prioritize alarms as well as predict life-threatening situations in the ICU. Our research uses clinical ICU patient data to develop, train, and test various ESN architectures for prediction tasks, and establishes the benefits of using ESN architecture designs for predicting ECG and ABP waveforms. Waveform and alarm classification and prediction is very important and we hope that this work will be helpful in providing additional insight into this problem.

2. Background

The recent release of clinical ICU patient data Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC II) makes it possible to develop better models and support tools to aid medical workers in understanding and filtering the abundance of information and alarms¹. This publicly available database has already been analyzed and used in several ways, for example to develop decision support systems to better categorize and classify mortality rates in the ICU^{6, 11-14}. A study previously used this data and an expert review panel to reclassify five common ICU alarms into true alarms and false alarms. These authors then developed an algorithm that classified alarms based on both electrocardiogram (ECG) and arterial blood pressure (ABP) waveforms immediately prior to the machine generated alarms. When tested, the algorithm suppressed approximately 59.7% of the false alarms and 0% of true alarms (except for true ventricular tachycardia alarms which was reduced by 9.4%)¹⁵. While their algorithm is focused on classification, this project aims to complement their work by developing a neural network that can learn to predict

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE NOV 2014		2. REPORT TYPE		3. DATES COVERED 00-00-2014 to 00-00-2014	
4. TITLE AND SUBTITLE Predicting Electrocardiogram and Arterial Blood Pressure Waveforms with Different Echo State Network Architectures				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory ,Washington,DC,20375				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES AMIA Annual Symposium, pp. 544-553, 15-19 Nov 2014, Washington, DC					
14. ABSTRACT Alarm fatigue caused by false alarms and alerts is an extremely important issue for the medical staff in Intensive Care Units. The ability to predict electrocardiogram and arterial blood pressure waveforms can potentially help the staff and hospital systems better classify a patient's waveforms and subsequent alarms. This paper explores the use of Echo State Networks, a specific type of neural network for mining, understanding, and predicting electrocardiogram and arterial blood pressure waveforms. Several network architectures are designed and evaluated. The results show the utility of these echo state networks, particularly ones with larger integrated reservoirs, for predicting electrocardiogram waveforms and the adaptability of such models across individuals. The work presented here offers a unique approach for understanding and predicting a patient's waveforms in order to potentially improve alarm generation. We conclude with a brief discussion of future extensions of this research.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 10	19a. NAME OF RESPONSIBLE PERSON
a REPORT unclassified	b ABSTRACT unclassified	c THIS PAGE unclassified			

the more continuous waveforms. An effective predictive model can help medical staff better anticipate a patient's condition, which includes the occurrence of alarms and false alarms.

It is difficult to predict time series data, especially with chaotic waveforms such as ECG and ABP. Previous studies that have focused on predicting ECG and ABP waveforms have used structures or approaches that simplified the waveforms, usually with higher order measures¹⁵⁻¹⁹. While there have been previous studies modeling and predicting complex time series, only a few studies have explored the relationship between ECG and ABP waveforms during classification or prediction tasks^{15, 20-24}. Our work explores the use of ESNs, a type of recurrent neural network, to explore this ECG-ABP relationship further. Neural networks have been shown to be good at predicting time series data, especially in situations where building proper heuristic models is difficult²⁵⁻²⁷. ESNs were chosen for this research because they have been previously shown to accurately predict chaotic time series without the need to train the specific internal representations of the system^{9, 10}. This computational advantage makes ESNs very attractive for predicting ECG and ABP time series which are both chaotic and difficult to learn.

2.1 Echo State Network

An Echo State Network (ESN), Figure 1, is an example of a recurrent neural network capable of modeling and predicting non-linear behaviors^{9, 10}. A typical ESN has four sets of unique weights: W_{in_hidden} , W_{hidden} , W_{hidden_out} , and W_{in_out} . W_{in_hidden} are randomly assigned fully connected weights between the input node to the reservoir nodes. A distinct property of ESNs is the sparsely connected, randomly assigned weights between the reservoir nodes, W_{hidden} . These sparsely connected nodes allow for pockets of local resonances, or echoes, to develop, which together can model complex waveforms^{9, 10}. Furthermore, W_{hidden_out} are randomly assigned fully connected weights from the reservoir nodes to the output node which contributes to the learning of the network from the teacher (or training) signal. Unlike more typical recurrent neural networks, ESNs do not have connections from the output nodes back to the hidden nodes; this greatly reduces the model complexity and convergence time for ESNs¹⁰. Lastly, W_{in_out} are randomly assigned weights from the input node to the output node.

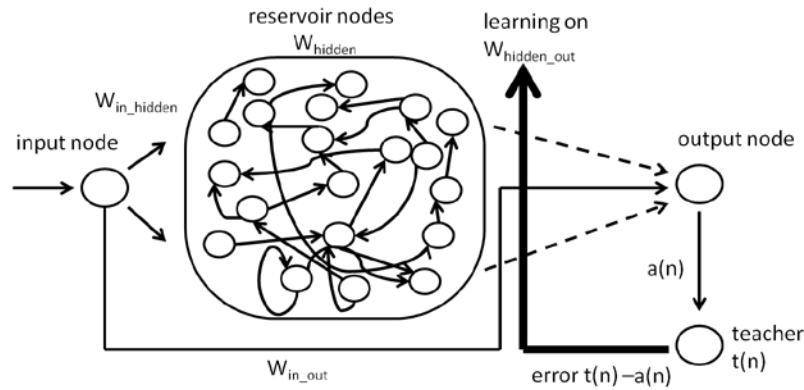


Figure 1: ESN architecture with learning on W_{hidden_out} (dashed arrows).

3. Methods

We comparatively evaluate the performance of three different types of ESN architectures at predicting two related physiological waveforms. The following sections first describe the data and preprocessing of the data. Next we discuss the analysis to identify reasonable ranges for key model parameters (reservoir size, activation rule, and learning rates). We then describe the three ESN architecture designs and our evaluation criteria.

3.1 Data Source and Pre-Processing

The data used for this project comes from the MIMIC II database which is publicly available^{11, 14, 28}. The complete database currently contains data from approximately 33,000 de-identified patients collected over 7 years (beginning in 2001) from Boston's Beth Israel Deaconess Medical Centers. It combines both clinical and physiological data. The adult patients range in age from 18 to over 90 years old (mean 68 years), and were collected from 48 medical, surgical, and coronary intensive care beds. Each patient record typically contains data from two electrocardiogram (ECG) leads, arterial blood pressure (ABP) and pulmonary arterial pressure (PAP) stored at 125Hz over time intervals that can range between a few hours to a few days. The ECG was originally sampled at 500Hz but was compressed to 125Hz while still preserving the peaks¹⁵. The resulting database is quite large (over 3TB). Because

we were interested in evaluating ESNs for individual patients, we focused on data from the ECG II (ECG lead II) and ABP readings from six randomly selected patients. ECG II data was selected because it appeared to be more available from a cursory look at the patient records. ABP was selected because of its relationship to ECG in the interpretation and classification of alarms¹⁵.

Although the data is publicly available, a specialized WaveForm DataBase (WFDB) software package was required to download, interpret, and format the data²⁹. Cygwin was used to connect directly to the server to download and format the data. The downloaded data was converted to comma separable version files which were compatible with Matlab. Some basic preprocessing was needed to make the magnitudes for the two waveforms comparable. A simple smoothing function, that averaged the data points in a 5-time step moving window, was applied to both the ECG and ABP data. This window size was effective at smoothing the waveform while maintaining the important features of the waves. The ABP data was also normalized to fall within the values 0 and 1. The ECG was vertically shifted up by the minimal value so it would be within the same range as the ABP data. These transformations were necessary in order to make the two waveforms similar in magnitude while maintaining unique features to allow for comparison.

3.2 Defining Baseline Model

An Echo State Network (ESN), Figure 1, was first built in Matlab based on previous work^{9, 10}. Weights were initially assigned random values between -0.5 and 0.5. Similar to previous work, approximately 20% of the possible connections in the reservoir have non-zero weights and are scaled with a spectral radius of 0.98, using:

$$W_{hidden} = \frac{\alpha W'_{hidden}}{|\lambda_{max}|}$$

where α is a scaling factor, W'_{hidden} is the weight matrix for the reservoir prior to transformation, and $|\lambda_{max}|$ is the maximum eigenvalue of W'_{hidden} , i.e., the spectral radius¹⁰.

The activation for the hidden nodes, A_{hidden} , and output nodes, A_{out} , is¹⁰:

$$A_{hidden}(t) = \tanh(W_{in_hidden}A_{in}(t) + W_{hidden}A_{hidden}(t-1))$$

$$A_{out}(t) = W_{in_out}A_{in}(t) + W_{hidden_out}A_{hidden}(t)$$

The learning rule and training only applied to the connections between the hidden nodes and the output nodes; all other weights remained unchanged through initialization, training, and testing. Although various learning techniques to train these weights were tried, such as linear regression and simplified error back propagation, a simple delta learning rule that incrementally changed the weights based on the product of the learning rate and the training error was shown to be both effective and fast. To prevent excessive oscillations in weights, a minimal error threshold was applied such that weights would not change if the absolute value of the error was less than 0.0001 (determined empirically).

To help validate this initial implementation, it was first tested by training it to model a simulated sine wave. The data was divided into training data (2,000 time-steps) and testing data (1,000 time-steps). The ESN, with 600 reservoir nodes and a learning rate of 0.0001, was initialized by passing the simulated sine wave through the reservoir once to let the internal system transients dissipate. Next, the training data was introduced to the network and the hidden-to-output weights were allowed to learn. The goal was to develop a model that could predict a waveform; hence the teacher signal to be predicted was the input signal 100 time-steps to the right (i.e., 100 time-steps in the future). This trivial example demonstrated the basic workings of this network. Mean Square Error (MSE) was used to evaluate this and subsequent test predictions:

$$MSE = \frac{1}{N} \sum_{n=1}^N (t(n) - a(n))^2$$

where N is the total number of time-steps in the test prediction, $t(n)$ is the actual teacher value at time-step n , and $a(n)$ is the output predicted value at time-step n . This trivial but useful prediction demonstration resulted in a low MSE test of 0.025.

3.3 ESN Architecture Designs

Three ESN architectures were designed to predict ECG and ABP waveforms. Three architectures were chosen to explore how coupling and integrating related waveforms effects overall predictive performance, Figure 2. The first

ESN consisted of two independent reservoirs (ESN1) which served as a control for the experiment. This architecture consisted of two ESNs time-synchronized and running in parallel with no connections between the networks. These reservoir sizes were determined experimentally as discussed later. The second coupled reservoir architecture (ESN2) was similar to the first architecture. However, the nodes in both reservoirs were connected to each of the output nodes. The learning of the hidden to output weights were specific to the error generated by the corresponding waveform. For example, only the error between the predicted and actual ECG waveforms was used to update the W_{hidden_out} connection to the ECG output node. Both reservoirs were initialized synchronously with their corresponding waveform. It was interesting to investigate if the predictions of one reservoir could benefit from the other reservoir with this architecture, and whether two separate reservoirs would make the overall network more robust to limitations associated with the randomly assigned weights.

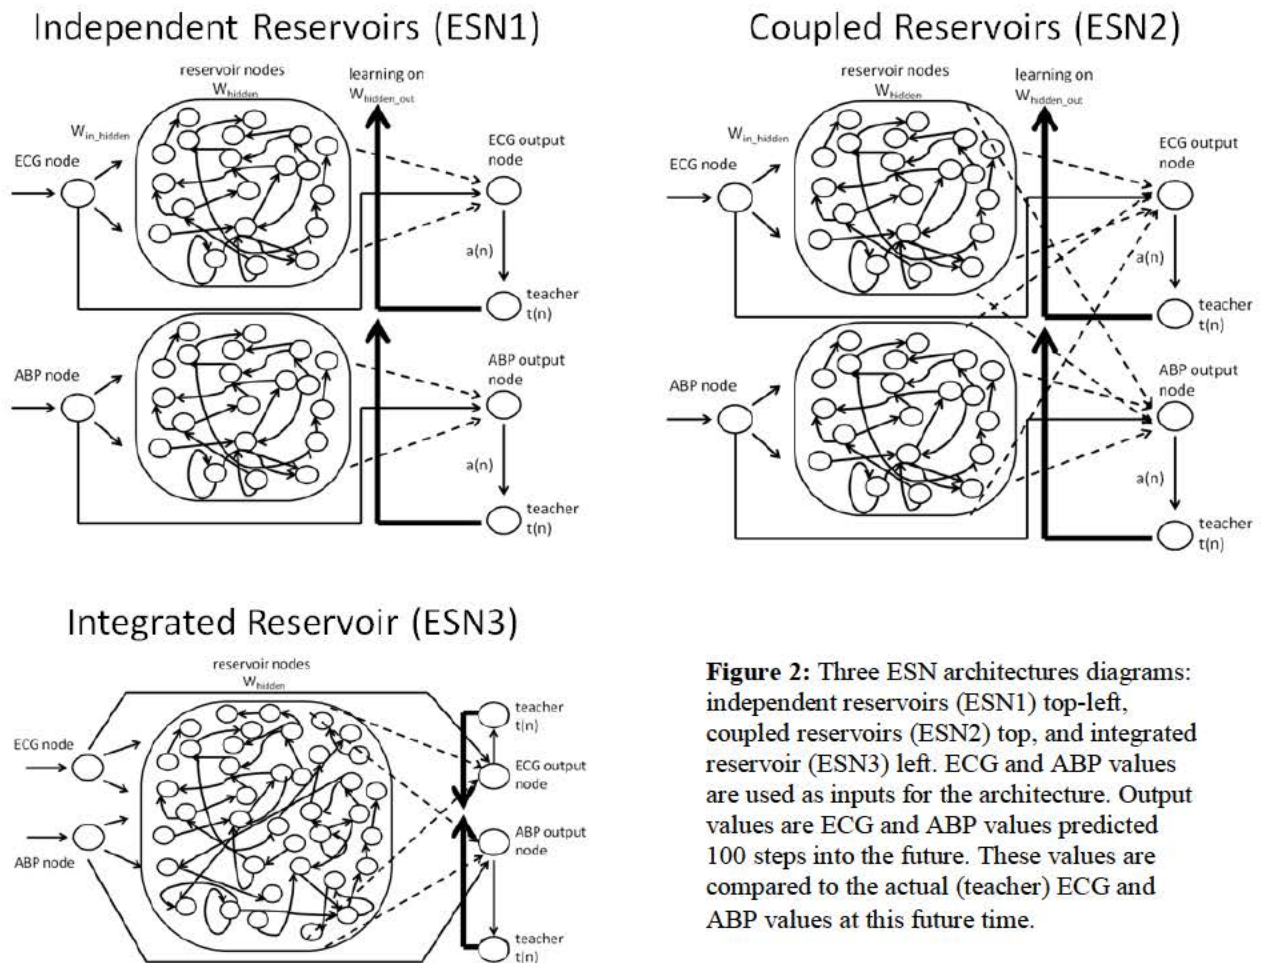


Figure 2: Three ESN architectures diagrams: independent reservoirs (ESN1) top-left, coupled reservoirs (ESN2) top, and integrated reservoir (ESN3) left. ECG and ABP values are used as inputs for the architecture. Output values are ECG and ABP values predicted 100 steps into the future. These values are compared to the actual (teacher) ECG and ABP values at this future time.

The third architecture consisted of one large integrated reservoir network (ESN3) with two input and two output nodes. This architecture was chosen to investigate how single reservoir systems compared to multiple reservoir systems when predicting multiple related waveforms. It was interesting to test if initializing and training one reservoir with two related waveforms could lead to better performance. To be consistent with the independent and coupled reservoirs, both the inputs and outputs were fully connected to the hidden nodes but inputs were only connected to their corresponding output waveform node. The other network parameters, such as spectral radius, W_{hidden} , sparsity, were the same with the independent and coupled reservoirs architectures. This ensured that the analysis would primarily focus on differences resulting from the network architecture. The integrated reservoir network was initialized with both waveforms, and similar to coupled reservoirs architecture, the fully connected

weights from the hidden to each output node were different and were trained based on the error associated with their corresponding predicted waveforms.

3.3.1 Determining Model Parameters

Experiments were first completed to identify some useful ESN parameter ranges for the ECG and ABP waveform data, especially reservoir sizes and learning rates. Identifying which parameters to set and the ranges of interest made the comparisons of different ESN architectures more appropriate. Two separate ESNs with similar architectures to the ESN mentioned above were used to predict ECG and ABP waveforms 100 time-steps into the future. These models were evaluated on MSE and maximum prediction error (max error). Different combinations of reservoir sizes and learning rates were tried because of their influence on how the waveforms were represented, decomposed, and learned by the system.

3.3.2 Methods for evaluation

The performance of these three architectures was assessed on MSE and the maximum prediction error (max error) between the performance and the actual waveform for both ECG and ABP data. A 10,000 time-step sample from patient record a41278 was used for this analysis. The data was divided into initializing (1-5,000), training (5,001-8,000), and testing segments (8,001-10,000). Initialization, training, and testing with all the networks followed the sample protocol. The Kolmogorov-Smirnov (KS) test was used to assess both the ECG and ABP MSE results for normality. This test is necessary to determine which statistical tests would be appropriate to use. Assuming normality, the ECG and ABP MSE and max error would be evaluated using one-way analysis of variance (ANOVA). This was used to determine if there were statistically significant differences between the ECG and ABP MSE and max error for the different architecture types (ESN1, ESN2, ESN3-900, ESN3-800, ESN3-700). Lastly, randomly sampled data from five individuals (a41325, a40416, a40076, a40432, and a41563) were used to evaluate the consistency and performance of the largest integrated reservoir network (ESN3-900).

4. Results

4.1 Model Parameters

To determine useful learning rates and reservoir sizes for the ESNs, a cursory assessment was first completed with the ECG data, varying the number of hidden nodes (100, 500, 750, and 1000) and the learning rates (0.01, 0.001, and 0.0001). The data was divided into initializing (1-5,000 time-steps), training (5,001-8,000 time-steps), and testing segments (8,001-10,000 time-steps). The data was tested 10 times for each combination. As shown in Table 1, reservoirs with nodes ranging between 100 and 750 and learning rates ranging from 0.001 and 0.0001 had on average better performance, prompting additional investigation as follows.

Table 1: ECG MSE test results (standard deviations)

MSE test	100	500	750	1000
0.01	3.78e42 (6.55e42)	7.83e7 (1.26e8)	4.1e13 (7.1e13)	4.6e76 (8.0e76)
0.001	0.026 (0.023)	18.39 (31.07)	1.42 (2.32)	1.6e6 (2.7e6)
0.0001	0.013 (0.0018)	0.02 (0.0029)	0.043 (0.031)	0.11 (0.14)

Reservoir size and learning rate ranges were further investigated with higher fidelity. The network ran ten more times with randomly initialized weights with new combinations of reservoir sizes (100, 200, 300, 400, 500, 600, and 700) and learning rates (0.001, 0.0005, and 0.0001). The results suggested that learning rates of 0.0001 and reservoir sizes of 500 or less than 300 tended to have better performance. A learning rate of 0.0001 and a reservoir size of 500 were used for the ECG components of the test ESN architectures. A reservoir size of 500 was chosen because it had slightly less variability in the results compared to sizes of 300 nodes or less. A similar analysis, conducted using ABP data, suggested a learning rate of 0.0001 and a reservoir size of 400 for the ABP components of the test ESN architectures. As a result, ESN1 and ESN2 both had two separate reservoirs with 500 and 400 nodes for the ECG and ABP waveforms respectively. Furthermore, we investigated the ESN3 architecture with three different reservoir sizes (700, 800, and 900 nodes) because a single reservoir system can perform differently based on its size.

4.2 Evaluating ESN Architecture Designs

The five ESN models were run thirty times with randomly initialized weights. Figures 3 and 4 show data from a training and testing run of ESN1 (independent reservoirs). The networks were evaluated on their prediction/test MSE and the maximum error values for both ECG and ABP waveforms. We found that the ESN architectures were able to predict ECG and ABP waveforms with varying levels of accuracy. We summarize the MSE results from both the ECG and ABP predictions in Figures 5 and 6 respectively. Outliers (results greater than three standard deviations from the mean) were most likely caused by poor randomly initialized weights and were removed, approximately five in each factor level, for the remainder of the analysis.

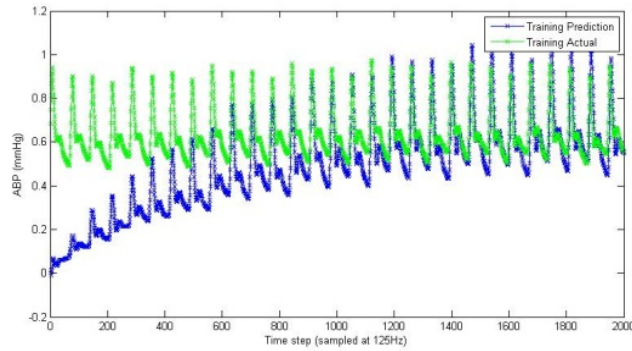


Figure 3: Sample ABP training run

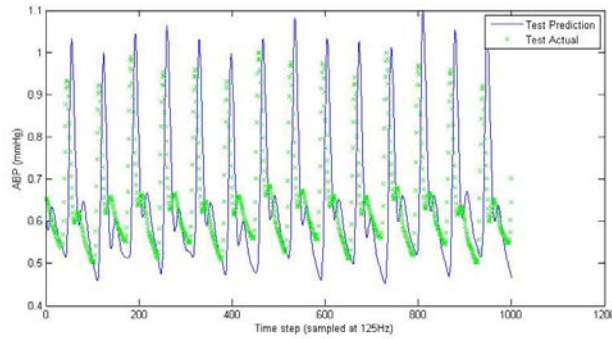


Figure 4: Sample ABP test prediction

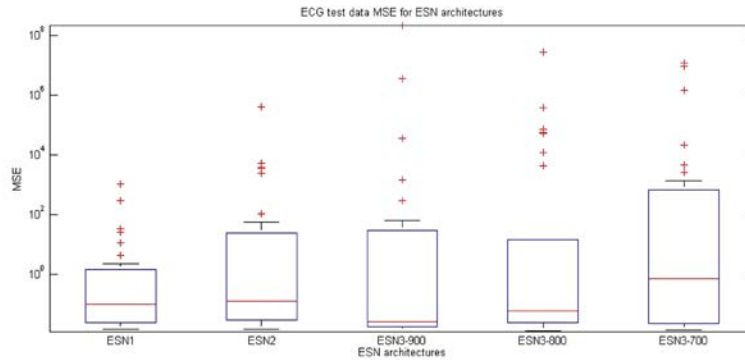


Figure 5: ECG MSE results for the different ESN architectures

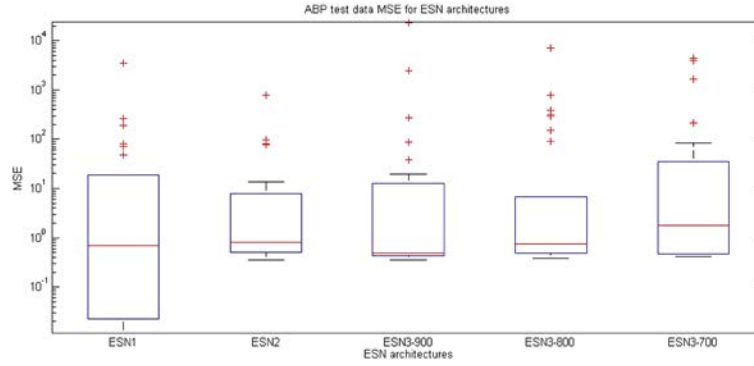


Figure 6: ABP MSE results for the different ESN architectures

Although the KS values for the ECG and ABP distributions were 0.5 and 0.53 respectively, parametric statistical tests could still be applicable considering the sample size. One-way ANOVAs were used to evaluate the MSE and max error of the predictions, Table 2. The factor levels were the different architectures (ESN1, ESN2, ESN3-900, ESN3-800, and ESN3-700).

Table 2: ANOVA results when evaluating ECG and ABP predictions using MSE and max error metrics

F-value (p-value)	ECG	ABP
MSE	1.48 (0.21)	1.19 (0.32)
Max Error	1.65 (0.17)	2.22 (0.071)

The results from the ABP max error analysis tended to be more meaningful (p-value = 0.071) compared to other metrics. The ABP max error results were mostly driven by the poor performance from ESN3-800 and ESN3-700 (integrated reservoir) as shown in Figure 7. We also noted that ESN3-900 (the largest integrated reservoir) tended to predict ECG waveforms slightly better than the independent and coupled reservoirs, Figure 7. However, ESN3's performance decreased with less hidden nodes. This may be due to the inability of the reservoir to correctly learn the two waveforms. In general, the independent reservoirs gave a much better prediction for the ABP waveform, although ESN3-900 had comparable performance in terms of ABP max error predictions. These results showed that there was no significant difference in the architectures at predicting two related waveforms together. There may also be increases in performance when combining waveforms in a single reservoir that is approximately similar in size to the ESN1 and ESN2. This showed the potential benefits of combining reservoirs to predict different but related waveforms.

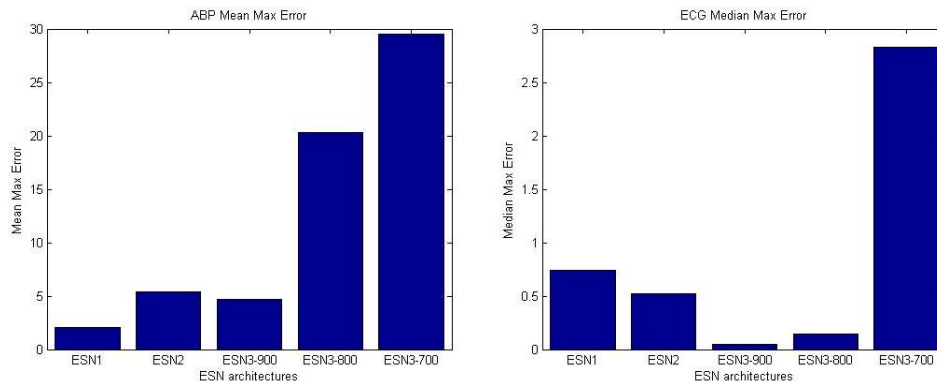


Figure 7: ABP mean max error (left) and ECG median max error (right).

4.3 Evaluating Across Individuals

Lastly, ESN3-900 (the largest integrated reservoir) was tested and compared using randomly sampled data from five individuals. ESN3-900 was chosen because its performance was similar to the basic independent reservoir architecture and it was faster to implement. The results showed that the ECG and ABP, Figure 8, performance across subjects were fairly similar, with the exception of patient 3. Though further work is needed, these results suggest that this architecture may be robust enough to be applied to different patients with little customization.

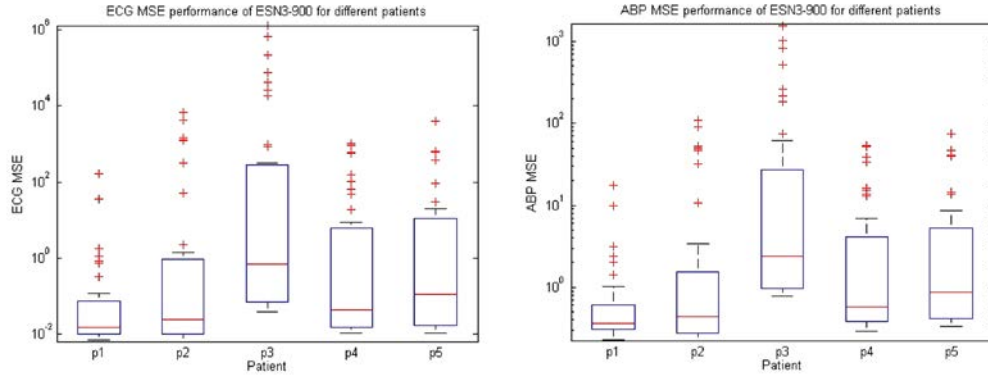


Figure 8: ECG MSE and ABP MSE results (left and right respectively) from five randomly selected patients

5. Discussion

The ability to predict and classify the ECG and ABP waveforms of patients is extremely important for the medical staff in Intensive Care Units. This paper approached this problem by demonstrating the effectiveness of different ESN architectures at predicting ECG and ABP waveforms. Although there was little significant difference in performance when predicting ECG and ABP waveforms between the different architectures, our results suggest that integrated reservoirs tended to have slightly better performance, especially when predicting the ECG waveform. Integrated reservoirs were also tested with different individuals with comparable results.

The ESN architectures in general tended to have much better performance at predicting ECG compared to ABP. This might be because the ECG waveforms have higher frequencies than the ABP waveforms. The loosely coupled subsystems in the reservoir might resonate or echo better at higher frequencies compared to lower frequencies. Furthermore, the ability of one larger reservoir to predict the waveforms was dramatically reduced as the size of the reservoir decreased. As the reservoir size decreased, the performance degradation was much more apparent with ABP than with ECG. This might also suggest that any coupling of ECG and ABP waveforms is biased toward the faster waveforms. This is an area that will require further investigation.

Furthermore, analysis of the largest integrated reservoir with five different individuals showed that this network architecture could be adaptable to different subjects. Without changing the parameters of the model, the largest integrated reservoir resulted in similar performance for four of the five randomly selected patients. Having a model that can be applied to different patients with little or no tailoring is attractive, and could be very beneficial in developing a predictive system for hospitals.

Although this research was exploratory, it does highlight some potential advantages of having one large reservoir for learning and predicting two related waveforms. The performance of the largest integrated reservoir was comparable to independent and coupled reservoirs, and was easier to implement. This analysis hints at interactions between how these waveforms are learned and stored in the reservoir. This could be investigated further and in more detail, perhaps starting with simpler, less chaotic waveforms. There are many questions to ask: for example, could initializing two reservoirs separately and then combining and reinitializing them lead to more robust internal subsystems in the reservoir? These and other queries can help further the understanding of Echo State Networks and make them more applicable for bioengineering and other applications.

There are limitations associated with this exploratory work, primarily concerning the limited predictive time-steps and the number of unique patients tested. This work aimed at exploring the application of ESNs for ECG and ABP waveforms. Although, these waveform predictions were limited to 100 time-steps, which may be too small to clinically trigger alarms, this work demonstrated the utility of this approach and could be refined in future work. It

would also be helpful to try different combinations of waveforms (not just the ECG II and ABP waveforms) and evaluate the models with more data from different patients. Furthermore, the parameters for each network (besides the learning rates and reservoir sizes) could be optimized. This would greatly increase the number of factors to control for, but a detailed factor level analysis of the network parameters would be very insightful. Expanding these models to predict higher level alarm states would be helpful. In addition, data transformations of the results might provide more normalized data. However, more research is needed to understand what these transformations mean intuitively for the results before applied.

A model that can predict ECG and ABP waveforms can naturally be extended to classify other waveforms and predict alarms. A model that can forecast the accuracy of ECG alarms or classify false alarms based on predicted waveforms might be extremely beneficial to medical staff, especially those in the ICU. There are several types of alarms in the ICU, each with unique waveform patterns, and this work can be extended to investigate the differences between predicted waveforms during a false alarm and a true alarm. Models that can extrapolate what a patient's waveforms will be like even a few seconds after an alarm can help medical staff and hospital systems better understand and classify alarms, with the ultimate goal of reducing false alarms, alarm fatigue, and improving patient care.

6. Conclusion

Alarm fatigue caused by bedside machines is a serious issue in Intensive Care Units. Part of this problem is the inability of these machines to accurately predict and classify a patient's ECG and ABP waveforms. In this study, we explored and demonstrated the ability of different Echo State Network architectures for predicting ECG and ABP waveforms with varying levels of accuracy. The most accurate predictions were generally by the largest integrated reservoir and the independent reservoirs architectures, which often had comparable results. Furthermore, results showed potential benefits for applying large integrated ESN reservoirs for different individuals. This paper also discussed limitations of this research, as well as suggestions for future work, and in particular the investigation of one versus two reservoir interactions and applications for predictive alarm classifications.

References

1. Mathews SC, Pronovost PJ. The need for systems integration in health care. *Journal of the American Medical Association*. 2011; 305(9): 934-5.
2. Graham KC, Cvach M. Monitor alarm fatigue: standardizing use of physiological monitors and decreasing nuisance alarms. *American Journal of Critical Care*. 2010; 19(1): 28-34.
3. Tsien CL, Fackler JC. Poor prognosis for existing monitors in the intensive care unit. *Critical care medicine*. 1997; 25(4): 614-619.
4. Cvach M. Monitor alarm fatigue: an integrative review. *Biomedical Instrumentation & Technology*. 2012; 46(4): 268-77.
5. Edelson, M. Safety First. *Hopkins Medicine*. 2013; 24-31.
6. Fuchs L, Chronaki CE, Park S, Novack V, Baumfeld Y, Scott D, et al. ICU admission characteristics and mortality rates among elderly and very elderly patients. *Intensive care medicine*. 2012; 38(10): 1654-61.
7. Konkani A, Oakley B, Bauld TJ. Reducing hospital noise: a review of medical device alarm management. *Biomedical Instrumentation & Technology*. 2012; 46(6): 478-87.
8. Jaeger H. Reservoir riddles: Suggestions for echo state network research. *Proceedings of the IEEE International Joint Conference of Neural Networks*; 2005; 3: 1460-2.
9. Jaeger H, Harald H. Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science*. 2004; 304(5667): 78-80.
10. Tong MH, Bickett AD, Christiansen EM, Cottrell GW. Learning grammatical structure with echo state networks. *Neural Networks*. 2007; 20(3): 424-32.
11. Saeed M, Villarroel M, Reisner AT, Clifford G, Lehman LW, Moody G, et al. Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): a public-access intensive care unit database. *Critical care medicine*. 2011; 39(5): 952.
12. Celi LA, Galvin S, Davidzon G, Lee J, Scott DJ, Mark RG. A Database-driven decision support system: customized mortality prediction. *Journal of personalized medicine*. 2012; 2(4): 138-48.
13. Celi LA, Mark RG, Lee J, Scott DJ, Panch T. Collective experience: a database-fuelled, inter-disciplinary team-led learning system. *Journal of computing science and engineering*. 2012; 6(1): 51.
14. Scott DJ, Lee J, Silva I, Park S, Moody GB, Celi LA, Mark RG. Accessing the public MIMIC-II intensive care relational database for clinical research. *BMC Medical Informatics and Decision Making*. 2013; 13(1): 9.

15. Aboukhalil A, Nielsen L, Saeed M, Mark RG, Clifford, GD. Reducing false alarm rates for critical arrhythmias using the arterial blood pressure waveform. *Journal of biomedical informatics*. 2008; 41(3): 442-51.
16. Fetics B, Nevo E, Chen CH, Kass DA. Parametric model derivation of transfer function for noninvasive estimation of aortic pressure by radial tonometry. *Biomedical Engineering*. 1999; 46(6): 698-706.
17. Keogh E, Lin J, Fu A. Hot sax: Efficiently finding the most unusual time series subsequence. *Proceedings of the IEEE International Joint Conference on Data Mining*; 2005;8.
18. Keogh E, Lin J, Lee SH, Van Herle H. Finding the most unusual time series subsequence: algorithms and applications. *Knowledge and Information Systems*. 2007; 11(1): 1-27.
19. Lonardi S, Lin J, Keogh E. Efficient discovery of unusual patterns in time series. *New Generation Computing*. 2006; 25(1): 61-93.
20. Zong W, Moody GB, Mark RG. Reduction of false arterial blood pressure alarms using signal quality assessment and relationships between the electrocardiogram and arterial blood pressure. *Medical and Biological Engineering and Computing*. 2004; 42(5): 698-706.
21. Saria S, Duchi A, Koller D. Discovering deformable motifs in continuous time series data. *International Joint Conference on Artificial Intelligence*. 2011; 22(1): 1465.
22. Williams C, Quinn J, McIntosh N. Factorial switching Kalman filters for condition monitoring in neonatal intensive care. *Neural Information Processing*. 2005; 18: 1513-1520.
23. Gather U, Imhoff M, Fried R. Graphical models for multivariate time series from intensive care monitoring. *Statistics in medicine*. 2002; 21(18): 2685-2701.
24. McSharry PE, Clifford GD, Tarassenko L, Smith L. A dynamical model for generating synthetic electrocardiogram signals. *IEEE Transactions on Biomedical Engineering*. 2003; 50(3): 289-294.
25. Kaastra I, Boyd M. Designing a neural network for forecasting financial and economic time series. *Neurocomputing*. 1996; 10(3): 215-36.
26. Maguire LP, Roche B, McGinnity TM, McDaid LJ. Predicting a chaotic time series using a fuzzy neural network. *Information Sciences*. 1998; 112(1): 125-36.
27. Zhang GP, Qi M. Neural network forecasting for seasonal and trend time series. *European Journal of Operational Research*. 2005; 160(2): 501-14.
28. MIMIC II [Internet]. MIT (MA): Multiparameter Intelligent Monitoring in Intensive Care [cited 2013 March 1]. Available from: <http://physionet.org/mimic2/>
29. WFDB Software Package [Internet]. MIT (MA): PhysioNet WFDB Software Package [cited 2013 March 1]. Available from: <http://physionet.org/physiotools/wfdb.shtml>